arXiv:1407.1096v2 [math.OC] 1 Jul 2016

SECOND-ORDER ORTHANT-BASED METHODS WITH ENRICHED HESSIAN INFORMATION FOR SPARSE ℓ_1–OPTIMIZATION

J.C. DE LOS REYES[‡], E. LOAYZA[‡] AND P. MERINO[‡]

Abstract. We present a second order algorithm, based on orthantwise directions, for solving optimization problems involving the sparsity enhancing ℓ_1-norm. The main idea of our method consists in modifying the descent orthantwise directions by using second order information both of the regular term and (in weak sense) of the ℓ_1-norm. The weak second order information behind the ℓ_1-term is incorporated via a partial Huber regularization. One of the main features of our algorithm consists in a faster identification of the active set. We also prove that a reduced version of our method is equivalent to a semismooth Newton algorithm applied to the optimality condition, under a specific choice of the algorithm parameters. We present several computational experiments to show the efficiency of our approach compared to other state-of-the-art algorithms.

1. Introduction

Optimization problems involving sparse vectors have important applications in fields like image restoration, machine learning, data classification, among many others [14, 35, 22, 10]. One of the most common mechanisms for enhancing sparsity consists in the use of the ℓ_1–norm of the vector in the cost function. This approach leads indeed to sparse solutions, but the design of numerical algorithms becomes challenging due to the non-differentiability of the ℓ_1–norm. Most of the algorithms developed to solve optimization problems with ℓ_1–norm consider in fact only first-order information, primal and/or dual, which seems natural due to the presence of the nondifferentiable term (see e.g. [35] and the references therein).

More recently, some authors have investigated the use of second order information in cases where the objective function is composed by a regular function and the ℓ_1–norm of the design variable [2, 4, 6]. By using second order information of the regular part, faster algorithms have been obtained. In the pioneering work [2], a limited memory BFGS method was considered in connection with so-called *orthantwise* directions. In [4] and [6] these type of directions were considered in connection with Newton and semismooth Newton type updates, respectively. Active set strategies based on second order information have also been recently envisaged ([31, 34, 40]).

This paper targets the question whether some useful information can also be extracted from the special structure of the ℓ_1–norm in order to design a second-order method. The main idea of our approach consists of incorporating second-order information "hidden" in the ℓ_1–norm. Roughly speaking, although the ℓ_1–norm is not differentiable in a classical sense, it has two derivatives in a distributional sense. By

2010 *Mathematics Subject Classification.* 49M15, 65K05, 90C53, 49J20, 49K20.

[*]This research has been supported by SENESCYT Award PIC-015-INAMHI-001 *"Sistema de Pronóstico del Tiempo para todo el Territorio Ecuatoriano: Modelización Numérica y Asimilación de Datos"*, a joint project between the Research Center on Mathematical Modelling (MODEMAT) and the Instituto Nacional de Meteorología e Hidrología (INAMHI). Moreover, we acknowledge partial support of MATHAmSud project SOCDE "Sparse Optimal Control of Differential Equations".

using a partial Huber regularization, we are able to extract that information for the update of the second order matrix, while keeping the same orthanwise descent type directions. The resulting algorithm enables a fast identification of the active set and, therefore, a fast decrease of the cost function. Let us remark that due to the choice of directions, our method solves the original nondifferentiable problem and therefore differs from fully regularized approaches like the one proposed in [13].

Our contribution encompasses also a study of the relation of our method with respect to semismooth Newton methods (**SSN**). We show that, under a specific choice of the defining constants of the semismooth Newton algorithm, the updates of a reduced version of our method turn out to be equivalent to the (**SSN**) updates. As a consequence, important convergence results are inherited from the generalized Newton framework, in particular local superlinear convergence. In addition, on basis of this equivalence, an adaptive algorithmic choice of the regularization parameter is proposed.

One of our main motivations for this work is the solution of sparse PDE-constrained optimization problems. In this field, the motivation for considering sparsity lies in the fact that such structure allows to localize the action of the controls, since sparse control functions have small support [36, 17, 7, 23]. Our main concern in this respect is that, whereas semismooth Newton methods applied to the optimality systems of elliptic problems work well (see [36]), their application to time-dependent problems requires, in general, to solve the full optimality system at once, which is computationally costly.

The outline of our paper is as follows. In Section 2 the main idea of our approach is described and the resulting orthantwise enriched second-order algorithm is presented. The main theoretical properties of our algorithm are studied in Sections 3 and 4. In Section 5 we show that a reduced version of our method is equivalent to a semismooth Newton method under a specific choice of the defining parameters. Finally, in Section 6 we exhaustively compare the behavior of our algorithm with other recently proposed methods.

1.1. **Problem formulation.** Let $f : \mathbb{R}^m \to \mathbb{R}$ and β a positive real number. We are interested in the numerical solution of the unconstrained optimization problem

$$\textbf{(P)} \qquad \min_{x\in\mathbb{R}^m} \quad \varphi(x) := f(x) + \beta\|x\|_1,$$

where $\|\cdot\|_1$ corresponds to the standard ℓ_1–norm in $\mathbb{R}^m$.

There are some interesting special cases for the cost function f. We mention, for instance:

- $f(x) = \|Ax - y\|_2^2$, where A is a matrix in $\mathbb{R}^{n\times m}$. This sparse linear regression problem receives the name of LASSO [22, 37].
- $f(x) = \|A^{-1}x - y\|_2^2 + \frac{\alpha}{2}\|x\|_2^2$. This particular loss function appears in linear-quadratic PDE–constrained optimization problems after a discretization of the partial differential operator. The function also involves an additional ℓ_2 Tikhonov regularization term [36, 7].
- $f(x) = -\frac{1}{N}\sum_{j=1}^{N} \log \frac{\exp(x_{y_j}^T z_j)}{\sum_{i\in C}\exp(x_i^T z_j)}$. This function appears in voice recognition problems, where N is the number of samples used for the recognition, C denotes the set of all class labels, y_j the label associated to the training points j, z_j is the feature vector and x_i is the parameters' subvector of class label i. The function represents the normalized sum of the negative log likelihood of each data point being placed in the correct class [10, 24].

Let us point out that our algorithm can also be used in presence of general regular functions f, without requiring convexity. This is of importance, in particular, for PDE-constrained optimization problems. Following [8] we will require the following conditions.

Assumption 1.

(i) $f : \mathbb{R}^m \to \mathbb{R}$ *is continuously differentiable, with* ∇f *Lipschitz continuous for some constant* $L > 0$.
(ii) $\varphi = f + \beta\|\cdot\|_1$ *is coercive.*

It is easy to argue existence of an optimal solution $\bar{x} \in \mathbb{R}^m$ for problem (**P**) under Assumption 1. Moreover, it is well known that depending on the size of the parameter β, the solution $\bar{x}$ tends to be more or less sparse. In fact, if f is convex and $\beta \geq \|\nabla f(0)\|_\infty$, then $\bar{x}$ is identically zero (see for instance [36]).

First-order optimality conditions for (**P**) can be obtained using standard tools, and can be stated as

$$0 \in \nabla f(\bar{x}) + \partial(\beta\|\bar{x}\|_1),$$

where $\partial\phi(x)$ denotes the subdifferential of the function ϕ at x. Moreover, the last inclusion is equivalent to the relations:

$$\begin{aligned}
&(1a) && 0 = && \nabla_i f(\bar{x}) + \beta && \text{for } i \in \bar{\mathcal{P}},\\
&(1b) && 0 = && \nabla_i f(\bar{x}) - \beta && \text{for } i \in \bar{\mathcal{N}},\\
&(1c) && 0 \in && [\nabla_i f(\bar{x}) - \beta, \nabla_i f(\bar{x}) + \beta] && \text{for } i \in \bar{\mathcal{A}},
\end{aligned}$$

where the index sets $\bar{\mathcal{P}}$, $\bar{\mathcal{N}}$ and $\bar{\mathcal{A}}$ are defined as

$$\bar{\mathcal{P}} = \{i : \bar{x}_i > 0\}, \quad \bar{\mathcal{N}} = \{i : \bar{x}_i < 0\}, \quad \text{and } \bar{\mathcal{A}} = \{i : \bar{x}_i = 0\}. \tag{2}$$

2. The orthant-wise enriched second-order method (OESOM)

In this section we present the main steps of the second-order algorithm we propose for solving (**P**). Since our method is based on the use of orthant directions, let us start by introducing them in a formal manner.

We define the *orthant directions* associated to a given vector $x \in \mathbb{R}^m$ as follows:

$$z_i(x) = \begin{cases} \operatorname{sign}(x_i) & \text{if } x_i \neq 0,\\ 1 & \text{if } x_i = 0 \text{ and } \nabla_i f(x) < -\beta,\\ -1 & \text{if } x_i = 0 \text{ and } \nabla_i f(x) > \beta,\\ 0 & \text{otherwise,} \end{cases} \tag{3}$$

for $i \in I := \{1, \ldots, m\}$, the index set, and where

$$\operatorname{sign}(x_i) := \begin{cases} 1 & \text{if } x_i > 0,\\ 0 & \text{if } x_i = 0,\\ -1 & \text{if } x_i < 0. \end{cases}$$

These directions actually correspond to the minimum norm subgradient element [35, Ch. 11].

A characterization of the orthant defined by $z(x)$ is given by:

$$\Omega := \{d \colon \operatorname{sign}(d) = \operatorname{sign}(z(x))\}. \tag{4}$$

Following [4] we consider the direction

$$\widetilde{\nabla}_i\varphi(x) = \begin{cases} \nabla_i f(x) + \beta\text{sign}(x_i) & \text{if } x_i \neq 0, \\ \nabla_i f(x) + \beta & \text{if } x_i = 0 \text{ and } \nabla_i f(x) < -\beta, \\ \nabla_i f(x) - \beta & \text{if } x_i = 0 \text{ and } \nabla_i f(x) > \beta, \\ 0 & \text{otherwise,} \end{cases} \tag{5}$$

which in the following will be called pseudo–gradient. It is worth notice that the pseudo–gradient $\widetilde{\nabla}\varphi(x)$ also belongs to $\nabla f(x) + \partial(\|x\|_1)$.

To ensure that the iterates remain in the orthant Ω, an additional projection step is required. The corresponding orthogonal projection is defined as:

$$\mathcal{P}(y)_i = \begin{cases} y_i & \text{if } \text{sign}(y_i) = \text{sign}(z_i(x)), \\ 0 & \text{otherwise.} \end{cases} \tag{6}$$

The use of this type of directions in combination with second-order updates was originally proposed in [2]. The resulting orthant-wise quasi-Newton matrix was constructed with the solely contribution of the regular part. In [4], instead of using the whole Hessian-like matrix for computing the direction, an additional Newton subspace correction step is performed. The efficiency of both methods was experimentally compared in [4].

As mentioned in the introduction, although the ℓ_1-norm is not differentiable in a classical sense, it is twice differentiable in a distributional sense (see, e.g., [9]). The second distributional derivative is given by Dirac's delta function:

$$\delta(x) = \begin{cases} +\infty & \text{if } x = 0, \\ 0 & \text{otherwise.} \end{cases}$$

Since this weak derivative lives only at a single point, it is usually dismissed. In our case, however, we are precisely interested in getting a large number of zero entries in the solution vector and, therefore, this information may become valuable.

To make this second-order information usable, let us consider a Huber regularization of the ℓ_1-norm given by $\|x\|_{1,\gamma} := \sum_{i=1}^m h_\gamma(x_i)$, where

$$h_\gamma(x_i) = \begin{cases} \gamma\frac{x_i^2}{2} & \text{if } |x_i| \leq \frac{1}{\gamma}, \\ |x_i| - \frac{1}{2\gamma} & \text{if } |x_i| > \frac{1}{\gamma}, \end{cases} \tag{7}$$

for $\gamma > 0$. The first partial derivatives are then given by

$$\nabla_i \|x\|_{1,\gamma} = \frac{\gamma x_i}{\max(1, \gamma|x_i|)}, \ i = 1, \ldots, m,$$

i.e., the gradient of the Huber regularization is a vector whose components are described by the formula above.

Since the first derivative is a semismooth function (see Section 5 below), it has also possible to compute a generalized second derivative. The generalized Hessian is a diagonal matrix with entries given by:

i) If $\gamma|x_i| \leq 1$, then

$$\Gamma_{ii} = \gamma.$$

ii) If $\gamma|x_i| > 1$, then

$$\Gamma_{ii} = \frac{\gamma}{\gamma|x_i|} - \frac{\gamma^2 x_i^2}{\gamma^2 x_i^2 |x_i|} = 0.$$

or, in a closed form, by

$$\Gamma_{ii} = \begin{cases} \gamma & \text{if } \gamma|x_i| \leq 1, \\ 0 & \text{elsewhere.} \end{cases} \tag{8}$$

By including this matrix in the second order system, together with the orthant directions, we obtain the following enriched system:

$$\left(B^k + \beta\Gamma^k\right) d^k = -\widetilde{\nabla}\varphi(x^k), \tag{9}$$

where B^k stands either for the Hessian of f or a symmetric positive definite quasi-Newton approximation of it (e.g. the BFGS matrix).

Remark 1. *Although several alternatives can be used, we consider a BFGS matrix B^k in our algorithm. The BFGS update in our case is constructed according to the well known formula:*

$$B^{k+1} = B^k - \frac{B^k\delta^k{\delta^k}^\top B^k}{{\delta^k}^\top B^k\delta^k} + \frac{y^k{y^k}^\top}{{y^k}^\top\delta^k}, \tag{10}$$

where $y^k = \nabla f(x^{k+1}) - \nabla f(x^k)$ and $\delta^k = x^{k+1} - x^k$. It should be noticed that there is a slightly difference with the classical BFGS method, since in our case the update of x^{k+1} varies from the classic BFGS method. This, however, does not affect the main properties of B_k, which is positive definite by construction.

Similarly to [2, 4], we consider the projected line-search rule

$$\varphi[\mathcal{P}(x^k + s_k d^k)] \leq \varphi(x^k) + \widetilde{\nabla}\varphi(x^k)^T[\mathcal{P}(x^k + s_k d^k) - x^k], \tag{11}$$

used in a backtracking procedure for choosing s_k.

Assumption 2. *There exist feasible line-search steps s_k, computed according to the line–search rule* (11)*, such that, for all $k \in \mathbb{N}$,*

$$\hat{s} \leq s_k \leq 1, \qquad \textit{for some } \hat{s} > 0. \tag{12}$$

The resulting algorithm is then given through the following steps.

Algorithm 1 Orthantwise Enriched Second Order Method (**OESOM**)

1: Initialize x^0 and B^0.
2: **repeat**
3: Compute the matrix Γ^k using (8) .
4: Compute the descent direction $\widetilde{\nabla}\varphi(x^k)$ using (5).
5: Compute d^k by solving the linear system (9).
6: Compute
$$x^{k+1} = \mathcal{P}(x^k + s_k d^k),$$
with the line–search step s_k computed by (11).
7: Update the matrix B^k.
8: $k \leftarrow k + 1$.
9: **until** stopping criteria is satisfied

3. Convergence Analysis

In this section we focus on the properties of the algorithm presented above. We verify that the directions of (**OESOM**) are in fact descent directions and prove convergence of the iterates.

Assumption 3. *Let B^k be a symmetric approximation of the Hessian of f in the k–th iteration. There exist positive constants $\hat{c}$ and $\hat{C}$ such that the following relation is satisfied:*

$$\hat{c}\|d\|_2^2 \leq d^\top B^k d \leq \hat{C}\|d\|_2^2, \qquad \text{for all } d \in \mathbb{R}^m. \tag{13}$$

From the assumptions on f and B^k, the second-order matrices B^k is bounded independent of k so that (9) implies that $\|\widetilde{\nabla}\varphi(x^k)\|_2 \leq c\|d^k\|$ for some constant c. In addition, from (9) we observe that for every i such that $|x_i^k| \leq \frac{1}{\gamma}$, we have

$$\sum_{\ell=1}^{m}(B_{\ell j}^k + \beta\Gamma_{\ell j})d_j = (B_{jj}^k + \beta\gamma)d_j + \sum_{\ell\neq j}^{m} B_{\ell j}d_j = -\widetilde{\nabla}_i\varphi(x^k),$$

which leads to the bound:

$$|d_j^k| = \left|\frac{1}{B_{jj}^k + \beta\gamma}\left(\sum_{\ell\neq j} B_{j\ell}^k d_\ell^k + \widetilde{\nabla}_j\varphi(x^k)\right)\right| \leq \frac{C}{\gamma}\|d\|_2, \tag{14}$$

for those j such that $|x_j^k| \leq 1/\gamma$.

At the k–th step of the algorithm, we define the strong active set by

$$\mathcal{S}_k := \{i\colon z_i^k = 0\}. \tag{15}$$

We recall that z_i^k depends on x^k. Hereafter, however, in order to simplify the notation, we avoid writing this dependence explicitely. From the definition of $\mathcal{S}_k$ it follows that, if $i \in \mathcal{S}_k$ then $x_i^k = 0$. The reciprocal is not always true. Moreover, we define the following index set of components of the current solution that remain in the orthant defined by z^k:

$$\mathcal{H}^k := \{i\colon \operatorname{sign}(x_i^k + s_k d_i^k) = \operatorname{sign}(z_i^k)\}. \tag{16}$$

With this definition we can express the $(k+1)$ iterate with help of the following update formula:

$$x_i^{k+1} = \mathcal{P}[x_i^k + s_k d_i^k] = x_i^k + \tilde{d}_i^k,$$

where

$$\tilde{d}_i^k = \begin{cases} s_k d_i^k & \text{if } i \in \mathcal{H}^k, \\ -x_i^k & \text{elsewhere.} \end{cases} \tag{17}$$

Lemma 1. *The orthant direction z^k defined in* (3) *and the direction $\tilde{d}^k$ satisfy*

$$(z^k)^\top \tilde{d}^k = \|x^k + \tilde{d}^k\|_1 - \|x^k\|_1. \tag{18}$$

Proof. By definition (17) we obtain that

$$\begin{aligned}\|x^k+\tilde{d}^k\|_1-\|x^k\|_1-(z^k)^\top\tilde{d}^k&=\sum_i\left(|x_i^k+\tilde{d}_i^k|-|x_i^k|\right)-\sum_i z_i^k\tilde{d}_i^k\\&=\sum_{i\in\mathcal{H}^k}\left(|x_i^k+s_kd_i^k|-|x_i^k|-z_i^ks_kd_i^k\right)-\sum_{i\notin\mathcal{H}^k}(|x_i^k|-z_i^kx_i^k)\\&=\sum_{i\in\mathcal{H}^k}\left(\mathrm{sign}(z_i^k)(x_i^k+s_kd_i^k)-|x_i^k|-z_i^ks_kd_i^k\right)\\&=\sum_{i\in\mathcal{H}^k}\left(z_i^k(x_i+s_kd_i^k)-|x_i^k|-z_i^ks_kd_i^k\right)=\sum_{i\in\mathcal{H}^k}z_i^kx_i^k-|x_i^k|=0\end{aligned}$$

■

Remark 2. *If f is twice continuously differentiable with a positive definite and Lipschitz continuous Hessian, then Assumption 3 is satisfied. The sequence of matrices $\{B^k\}$ approximating the Hessian can be constructed, for instance, with the BFGS method, which leads to symmetric and positive definite matrices, if the curvature condition is satisfied. We point out that the next results are valid for any positive definite approximation of the Hessian engendered by some quasi–Newton method satisfying Assumption 3.*

Proposition 1. *Let x^k be the k–th iterate of the algorithm and z^k its associated orthant direction. Let γ be sufficiently large and s_k sufficiently small such that*

$$\mathrm{sign}(x_i^k+s_kd_i^k)=\mathrm{sign}(x_i^k),\ \textit{for all } i:\ |x_i^k|\geq\frac{1}{\gamma}. \tag{19}$$

Then, under Assumptions 2 and 3 we have that

$$\widetilde{\nabla}\varphi(x^k)^\top\tilde{d}^k<-\left(\frac{\hat{c}\hat{s}}{2}-s_k\,\mathcal{O}\left(\gamma^{-2}\right)\right)\|d^k\|_2^2-\hat{c}\|\tilde{d}^k\|_2^2. \tag{20}$$

In particular, if s is bounded and γ is sufficiently large there exists a constant $\mu>0$, independent of k, such that

$$\widetilde{\nabla}\varphi(x^k)^\top\tilde{d}^k\leq-\mu\left(\|d^k\|_2^2+\|\tilde{d}^k\|_2^2\right). \tag{21}$$

Proof. Let us simplify our presentation by dropping off the superindex k on the variables generated at the k–th iteration. In this way d and $\tilde{d}$ denote the vectors defined by (9) and (17), respectively. In addition, let us denote by $\hat{d}$ the vector whose components are given by

$$\hat{d}_i=\begin{cases}x_i+sd_i, & \text{if } i\notin\mathcal{H}^k,\\ 0, & \text{otherwise,}\end{cases} \tag{22}$$

with $\mathcal{H}^k$ defined by (16). Therefore, we can write $\tilde{d}=sd-\hat{d}$.

From the positive definiteness of B we obtain the inequality:

$$\hat{c}\|\tilde{d}\|_2^2\leq\tilde{d}^\top B\tilde{d}=(sd-\hat{d})^\top B(sd-\hat{d})=s^2d^\top Bd-2sd^\top B\hat{d}+\hat{d}^\top B\hat{d}. \tag{23}$$

Taking the second term to the left–hand side and dividing by $2s$, we get

$$d^\top B\hat{d}\leq\frac{s}{2}d^\top Bd+\frac{1}{2s}\hat{d}^\top B\hat{d}-\hat{c}\|\tilde{d}\|_2^2. \tag{24}$$

Now we proceed to analyze the quantity $\tilde{\nabla}\varphi(x)^\top \tilde{d}$. For this, we take into account equation (9) and obtain:

$$\begin{aligned}\tilde{\nabla}\varphi(x)^\top \tilde{d} &= -d^\top (B+\beta\Gamma)\,\tilde{d}\\ &= -d^\top (B+\beta\Gamma)\,(sd-\hat{d})\\ &= -sd^\top (B+\beta\Gamma)d + d^\top B\,\hat{d} + \beta d^\top \Gamma\,\hat{d}. \end{aligned} \tag{25}$$

Applying (24) and Assumption 3, we get that

$$\begin{aligned}\tilde{\nabla}\varphi(x)^\top \tilde{d} &\le -sd^\top (B+\beta\Gamma)d + \frac{s}{2}d^\top Bd + \frac{1}{2s}\hat{d}^\top B\hat{d} + \beta d^\top\Gamma\,\hat{d} - \hat{c}\|\tilde{d}\|_2^2\\ &= -\frac{s}{2}d^\top Bd - s\beta\, d^\top\Gamma d + \frac{1}{2s}\hat{d}^\top B\hat{d} + \beta d^\top\Gamma\,\hat{d} - \hat{c}\|\tilde{d}\|_2^2\\ &\le -\hat{c}\frac{s}{2}\|d\|_2^2 - s\beta\, d^\top\Gamma d + \frac{1}{2s}\hat{d}^\top B\hat{d} + \beta d^\top\Gamma\,\hat{d} - \hat{c}\|\tilde{d}\|_2^2. \end{aligned} \tag{26}$$

In the last inequality, the second term on the right–hand side is non-positive in view of positive semi-definiteness of Γ. Let us analyze the third term in (26). It is clear that this term is positive and, by definition, the components of $\hat{d}$ are those that change its sign with respect to the current orthant. Therefore,

$$|x_i + sd_i| \le s|d_i| \text{ for } i \notin \mathcal{H}^k, \tag{27}$$

since $\text{sign}(x_i) = \text{sign}(z_i)$ for $x_i \neq 0$.

Therefore, by the definition of $\hat{d}$ in (22) and Assumption 3, we have that

$$\begin{aligned}\frac{1}{2s}\hat{d}^\top B\hat{d} &\le \frac{\hat{C}}{2s}\|\hat{d}\|_2^2\\ &\le \frac{\hat{C}}{2s}\sum_{i\notin\mathcal{H}^k} s^2|d_i|^2\\ &= \frac{s}{2}\hat{C}\Big(\sum_{i:i\notin\mathcal{H}^k,|x_i|\le\frac{1}{\gamma}} |d_i|^2 + \sum_{i:i\notin\mathcal{H}^k,|x_i|>\frac{1}{\gamma}} |d_i|^2\Big). \end{aligned} \tag{28}$$

Now, by taking into account inequality (14) we get

$$\frac{1}{2s}\hat{d}^\top B\hat{d} \le \frac{s}{2}\hat{C}\sum_{i:i\notin\mathcal{H}^k,|x_i|>\frac{1}{\gamma}} |d_i|^2 + s\,\mathcal{O}\left(\frac{1}{\gamma^2}\right)\|d\|_2^2. \tag{29}$$

By similar arguments, thanks to (27) and since $\hat{d}_i = 0$ if $i \in \mathcal{H}_k$, we get that

$$\beta d^\top\Gamma\,\hat{d} \le \beta\gamma \sum_{i:i\notin\mathcal{H}^k,|x_i|\le\frac{1}{\gamma}} d_i(x_i+sd_i) \le s\beta\gamma \sum_{i:i\notin\mathcal{H}^k,|x_i|\le\frac{1}{\gamma}} |d_i|^2. \tag{30}$$

By plugging (29) and (30) in (26), we arrive at

$$\begin{aligned}\tilde{\nabla}\varphi(x)^\top \tilde{d} &\le -\hat{c}\frac{s}{2}\|d\|_2^2 + \frac{s}{2}\hat{C}\sum_{i:i\notin\mathcal{H}^k,|x_i|>\frac{1}{\gamma}} |d_i|^2 - \hat{c}\|\tilde{d}\|_2^2 + s\,\mathcal{O}\left(\gamma^{-2}\right)\|d\|_2^2\\ &\le -\left(\frac{\hat{c}\hat{s}}{2} - s\,\mathcal{O}\left(\gamma^{-2}\right)\right)\|d\|_2^2 - \hat{c}\|\tilde{d}\|_2^2, \end{aligned} \tag{31}$$

where the last inequality follows from Assumption 2 and taking into account that (19) implies that the set $\{i : i \notin \mathcal{H}^k, |x_i| > \frac{1}{\gamma}\}$ becomes empty. Our assertion is obtained by taking γ sufficiently large. ∎

Remark 3. *Condition* (19) *can be directly satisfied by choosing*

$$s_k < \min_{i \notin \mathcal{H}^k, |x_i^k| > 1/\gamma} \left| \frac{x_i^k}{d_i^k} \right|.$$

Theorem 1. *Let us assume that Assumption 3 and condition* (19) *hold true. If the direction* $\tilde{d}^k$ *defined by* (17) *is different from* 0*, then it satisfies*

$$\varphi(x^k + \tilde{d}^k) < \varphi(x^k), \tag{32}$$

for a sufficiently small step size s_k*.*

Proof. Using Lemma 1 and Proposition 1, and using a Taylor expansion of f at x^k we get

$$\begin{aligned}
\varphi(x^k + \tilde{d}^k) &= f(x^k + \tilde{d}^k) + \beta \|x^k + \tilde{d}^k\|_1 \\
&= f(x^k) + \nabla f(x^k)^\top \tilde{d}^k + o(\|\tilde{d}^k\|_2) + \beta \|x^k + \tilde{d}^k\|_1 \\
&= \varphi(x^k) + \nabla f(x^k)^\top \tilde{d}^k + \beta(\|x^k + \tilde{d}^k\|_1 - \|x^k\|_1) + o(\|\tilde{d}^k\|_2) \\
&= \varphi(x^k) + \nabla f(x^k)^\top \tilde{d}^k + \beta z^\top \tilde{d}^k + o(\|\tilde{d}^k\|_2) \\
&= \varphi(x^k) + (\nabla f(x^k) + \beta z)^\top \tilde{d}^k + o(\|\tilde{d}^k\|_2).
\end{aligned}$$

Considering (5) we get

$$\begin{aligned}
\varphi(x^k + \tilde{d}^k) &= \varphi(x^k) + \tilde{\nabla}\varphi(x^k)^\top \tilde{d}^k + \sum_{i:\, x_i = 0,\, |\nabla f(x_i)| \le \beta} \nabla_i f(x_i) \tilde{d}_i^k + o(\|\tilde{d}^k\|_2) \\
&\le \varphi(x^k) + \tilde{\nabla}\varphi(x^k)^\top \tilde{d}^k + \beta \sum_{i:\, x_i = 0,\, |\nabla f(x_i)| \le \beta} |\tilde{d}_i^k| + o(\|\tilde{d}^k\|_2). \qquad (33)
\end{aligned}$$

By definition of $\tilde{d}^k$, we have that $\tilde{d}_i^k = s_k d_i^k$ for $i \in \mathcal{H}^k$. If $i \notin \mathcal{H}^k$ and $x_i^k = 0$, then $\tilde{d}_i^k = -x_i^k = 0$. In the case that $i \notin \mathcal{H}^k$ and $x_i^k \neq 0$ we use again (27). Altogether, we get that $|\tilde{d}_i^k| \le s_k |d_i^k|$, for all i. Consequently, $o(\|\tilde{d}^k\|_2) = o(|s_k|)$ since d^k is bounded. Therefore, in view of Assumption 3, together with (33) and (21), we can estimate

$$\varphi(x^k + \tilde{d}^k) \le \varphi(x^k) - \mu \left(\|d^k\|_2^2 + \|\tilde{d}^k\|_2^2 \right) + \beta s_k \sum_{i:\, x_i = 0,\, |\nabla f(x_i)| \le \beta} |d_i| + o(s_k). \tag{34}$$

Finally, we notice that in the set of indexes i where $x_i^k = 0$ and $|\nabla f(x_i^k)| \le \beta$ we get that $z_i^k = 0$. Therefore, if $i \in \mathcal{H}_i^k$, then $\operatorname{sign}(x_i^k + s_k d_i^k) = \operatorname{sign}(s_k d_i^k) = \operatorname{sign}(z_i^k) = 0$. On the other hand, if $i \notin \mathcal{H}_i^k$ then $\tilde{d}_i^k = -x_i^k = 0$, because of the definition of $\tilde{d}_i^k$. By this analysis we conclude that the sum term in (34) vanishes and we get the estimate

$$\varphi(x^k + \tilde{d}^k) \le \varphi(x^k) - \mu \left(\|d^k\|_2^2 + \|\tilde{d}^k\|_2^2 \right) + o(s_k). \tag{35}$$

This shows that $\varphi(x^k + \tilde{d}^k) < \varphi(x^k)$ for s_k sufficiently small and proves that $\tilde{d}^k$ is a descent direction. ∎

The following result is proved in [8, Lemma 4.3] and will be used in our convergence analysis.

Lemma 2 (Chouzenoux et al. (2014))**.** *Let* $(u^k)_{k\in\mathbb{N}}$*,* $(g^k)_{k\in\mathbb{N}}$*,* $(\hat{g}^k)_{k\in\mathbb{N}}$ *and* $(\hat{\Delta}^k)_{k\in\mathbb{N}}$ *be sequences of nonnegative reals and let* $\theta \in]0, 1[$*. Assume that*

(i) for every $k \in \mathbb{N}$*,* $(u^k)^2 \le (g^k)^\theta \hat{\Delta}^k$*,*
(ii) $(\hat{\Delta}^k)_{k\in\mathbb{N}}$ *is summable,*

(iii) for every $k \in \mathbb{N}$, $g^{k+1} \leq (1-\alpha)g^k + \hat{g}^k$, *where* $\alpha \in]0,1]$, *and*
(iv) for every $k \geq k_0$, $(\hat{g}^k)^\theta \leq \rho u^k$, *where* $\rho > 0$ *and* $k_0 \in \mathbb{N}$.
Then $(u^k)_{k\in\mathbb{N}}$ *is a summable sequence.*

In the next result we will rely on the following form of the Kurdyka–Łojasiewicz property (see [8]): A function ϕ satisfies the Kurdyka–Łojasiewicz inequality if for every $\xi \in \mathbb{R}$ and for every bounded subset $E \subset \mathbb{R}^m$, there exist three constants $\kappa > 0$, $\zeta > 0$ and $\theta \in [0,1[$ such that for all $v \in \partial\phi(x)$ and every $x \in E$ such that $|\phi(x) - \xi| \leq \zeta$, it follows that

$$\kappa|\phi(x) - \xi|^\theta \leq \|v\|_2, \tag{36}$$

with the convention $0^0 = 0$.

Assumption 4. *The function* φ *defined in* (**P**) *satisfies the Kurdyka–Łojasiewicz property, i.e., there exist positive constants* κ, ζ *and* $\theta \in [0,1[$ *such that for all* $z \in \partial(\beta\|\cdot\|_1)(x)$ *and every* $x \in E$ *(with* $E \subset \mathbb{R}^m$ *a bounded set), it holds that*

$$\kappa|\varphi(x) - \xi|^\theta \leq \|\nabla f(x) + z\|_2, \tag{37}$$

for $\xi \in \mathbb{R}$ *such that* $|\varphi(x) - \xi| \leq \zeta$.

The next result is based on the ideas developed in the convergence analysis presented in [8, Theorem 4.1]. There, the Kurdyka–Łojasiewicz property is applied in order to show that the sequence $(\|x^{k+1} - x^k\|)_{k\in\mathbb{N}}$ fulfills a finite length property, which allows to conclude the convergence of $(x^k)_{k\in\mathbb{N}}$. In our case, the convergence analysis is analogous to the theory developed in [8] but there are some key differences. For convenience of the reader, we include a detailed proof which mimics some of the steps of [8, Theorem 4.1].

Theorem 2. *Let us suppose that Assumptions 1-4 are satisfied. Then* $(x^k)_{k\in\mathbb{N}}$ *generated by Algorithm 1 converges to point* $\bar{x}$ *such that* $0 \in \nabla f(\bar{x}) + \partial(\beta\|\bar{x}\|_1)$.

Proof. Since the restriction of φ to its domain is continuous, and by Assumption 1, φ is coercive, we have that the level set $\{x : \varphi(x) \leq \varphi(x^0)\}$ is a compact set. Moreover, by Theorem 1 we have the monotonicity property: $\varphi(x^{k+1}) < \varphi(x^k)$, for all $k \in \mathbb{N}$. In addition, since φ is bounded from below, we have that the sequence $(\varphi(x^k))_{k\in\mathbb{N}}$ converges to some limit ξ as $k \to \infty$.

On the other hand, let s_k be a step–size satisfying (35), Theorem 1, and Assumptions 2 and 3. We have that, for sufficiently large γ, there is a constant $\hat{\mu}$ such that

$$\hat{\mu}\|d^k\|_2^2 \leq \varphi(x^k) - \varphi(x^{k+1}), \tag{38}$$

and, therefore,

$$\|d^k\|_2^2 \leq \frac{1}{\hat{\mu}}\Big((\varphi(x^k) - \xi) - (\varphi(x^{k+1}) - \xi)\Big). \tag{39}$$

Note that (39) implies that $\|d^k\|_2 \to 0$ as $k \to \infty$. In order to estimate the right–hand side of (39), we will use the following inequality for convex differentiable functions $\psi : [0,+\infty[\to [0,+\infty[$:

$$\psi(u) - \psi(v) \leq \psi'(u)(u - v), \tag{40}$$

for all u and v in the interval $[0,+\infty[$. By taking $\psi(u) = u^{\frac{1}{1-\theta}}$, with $\theta \in [0,1[$, $u = (\varphi(x^k) - \xi)^{1-\theta}$ and $v = (\varphi(x^{k+1}) - \xi)^{1-\theta}$, (40) implies that

(41)
$$(\varphi(x^k)-\xi)-(\varphi(x^{k+1})-\xi)\leq\frac{1}{1-\theta}(\varphi(x^k)-\xi)^\theta\left[(\varphi(x^k)-\xi)^{1-\theta}-(\varphi(x^{k+1})-\xi)^{1-\theta}\right].$$

Let us denote $\Delta^k:=(\varphi(x^k)-\xi)^{1-\theta}-(\varphi(x^{k+1})-\xi)^{1-\theta}$. From (39) and (41) we get

$$\|d^k\|_2^2\leq\frac{1}{\hat{\mu}(1-\theta)}(\varphi(x^k)-\xi)^\theta\Delta^k. \tag{42}$$

By Assumption 4, φ satisfies the Kurdyka–Łojasiewicz property with constants κ, ζ and θ. Moreover, since $\varphi(x^k)\to\xi$, there exists k_0 such that for $k\geq k_0$ we have that $|\varphi(x^k)-\xi|\leq\zeta$ and therefore, by (37) we obtain

$$\kappa|\varphi(x^k)-\xi|^\theta\leq\|\nabla f(x^k)+z\|_2,\quad\forall\, z\in\partial\left(\beta\|\cdot\|_1\right)(x^k). \tag{43}$$

In particular, if we choose $\tilde{z}\in\partial\left(\beta\|\cdot\|_1\right)(x^k)$ as

$$\tilde{z}_i=\begin{cases}\beta\ \mathrm{sign}(x_i) & \text{if } x_i\neq 0,\\ \beta & \text{if } x_i=0 \text{ and } \nabla_i f(x^k)<-\beta,\\ -\beta & \text{if } x_i=0 \text{ and } \nabla_i f(x^k)>\beta,\\ -\nabla_i f(x^k) & \text{otherwise,}\end{cases} \tag{44}$$

from (43) we can estimate

$$\begin{aligned}\kappa|\varphi(x^k)-\xi|^\theta&\leq\|\nabla f(x^k)+\tilde{z}\|_2=\|\widetilde{\nabla}\varphi(x^k)\|_2,\\ &=\|(B^k+\beta\Gamma^k)d^k\|_2\leq\|B^k+\beta\Gamma^k\|\|d^k\|_2,\end{aligned} \tag{45}$$

and by our Assumption 3, it follows that

$$|\varphi(x^k)-\xi|^\theta\leq\frac{C}{\kappa}\|d^k\|_2, \tag{46}$$

for some positive constant C.

Now we use Lemma 2 to infer the summability of the sequence $(u^k)_{k\in\mathbb{N}}$ defined by $u^k=\|d^k\|$. For this purpose we define the nonnegative sequences $(g^k)_{k\in\mathbb{N}}$, $(\hat{g}^k)_{k\in\mathbb{N}}$ and $(\hat{\Delta}^k)_{k\in\mathbb{N}}$ by $g^k=\varphi(x^k)-\xi$, $\hat{g}^k=\varphi(x^{k+1})-\xi$ and $\hat{\Delta}^k=\dfrac{1}{\hat{\mu}(1-\theta)}\Delta^k$, respectively. Let us verify conditions (i)–(iv) in Lemma 2:

(i) is satisfied considering (42): $\|d^k\|_2^2\leq(g^k)^\theta\hat{\Delta}^k$.

(ii) is obtained following [8, Theorem 4.1]:

$$\begin{aligned}\sum_{k\geq k_0}^{\infty}\Delta^k&=\sum_{k\geq k_0}^{\infty}\left[(\varphi(x^k)-\xi)^{1-\theta}-(\varphi(x^{k+1})-\xi)^{1-\theta}\right]\\ &=(\varphi(x^{k_0})-\xi)^{1-\theta},\end{aligned}$$

which implies that $(\hat{\Delta}^k)_{k\in\mathbb{N}}$ is summable.

(iii) follows from the choice $g^{k+1}=\varphi(x^{k+1})-\xi=\hat{g}^k=(1-\alpha)g^k+\hat{g}^k$, with $\alpha=1$.

(iv) From Theorem 1 we have that $\varphi(x^{k+1})-\varphi(x^k)<0$ and

$$\begin{aligned}(\hat{g}^k)&=(\varphi(x^{k+1})-\xi)\\ &=\varphi(x^{k+1})-\varphi(x^k)+\varphi(x^k)-\xi\\ &<\varphi(x^k)-\xi.\end{aligned}$$

Therefore, by using (46), for $k \geq k_0$, we have

$$(\hat{g}^k)^\theta < (\varphi(x^k) - \xi)^\theta \leq \frac{C}{\kappa}\|d^k\|_2.$$

Consequently, conditions (i)–(iv) of Lemma 2 are satisfied. This allows us to conclude that $\|d^k\|$ is summable when $\theta > 0$. If $\theta = 0$, from the convention $0^0 = 0$ and since, for sufficiently large k, $\frac{C}{\kappa}\|d^k\| < 1$, (46) implies that $\varphi(x^k) = \xi$ (otherwise a contradiction is obtained). Replacing this last identity in (42), it follows that $\|d^k\| = 0$ and $\|d^k\|$ is summable.

From the summability of $\|d^k\|$ we conclude the summability of $\|\tilde{d}^k\|$ since $\|\tilde{d}^k\| \leq \|d^k\|$, for all $k \in \mathbb{N}$.

Our final step consists in verifying the convergence of $(x^k)_{k\in\mathbb{N}}$ to a critical point of φ. Since $\|x^{k+1} - x^k\|_2 = \|\tilde{d}^k\|_2 \to 0$ as $k \to \infty$, and $\|\tilde{d}^k\|$ is summable, then $(x^k)_{k\in\mathbb{N}}$ is a Cauchy sequence (then convergent). Let us denote its limit by $\bar{x}$.

Since $\widetilde{\nabla}\varphi(x^k) \in \nabla f(x^k) + \partial(\beta\|\cdot\|_1)(x^k)$ we have that the pair $(x^k, \widetilde{\nabla}\varphi(x^k)) \in \mathrm{Graph}(\nabla f + \partial(\beta\|\cdot\|_1))$. Proceeding as in (45), we deduce $\|\widetilde{\nabla}\varphi(x^k)\|_2 \leq \frac{C}{\kappa}\|d^k\|_2$, which by (39) yields

$$(x^k, \widetilde{\nabla}\varphi(x^k)) \to (\bar{x}, 0) \quad \text{as} \quad k \to +\infty,$$

which belongs to $\mathrm{Graph}(\nabla f + \partial(\beta\|\cdot\|_1))$ due to its closedness. Therefore we obtain that $0 \in \nabla f(\bar{x}) + \partial(\beta\|\cdot\|_1)(\bar{x})$, from which we finally get the result. ∎

4. Active set properties

In the following theorem, we show that the identification of the strongly active sets performed by (**OESOM**) is monotone in a neighbourhood of the solution $\bar{x}$. The size of the neighbourhood turns out to be strongly dependent on the regularization parameter γ. The monotonicity of the active sets is proved under a strict complementarity assumption.

Theorem 3 (Monotonicity of the strongly active set)**.** *Let us assume that the sequence $\{x^k\}_{k\in\mathbb{N}}$ generated by algorithm (***OESOM***) converges to the solution $\bar{x}$. If strict complementarity holds at x^k, i.e. $\{i : \nabla_i f(x^k) = \beta, x_i^k = 0\} = \emptyset$, then, for k sufficiently large,*

$$\mathcal{S}_k \subseteq \mathcal{S}_{k+1}. \tag{47}$$

Proof. Let us suppose that $z_i^k = 0$ for some index i. In order to prove (47) we must verify two properties: *i)* $x_i^{k+1} = 0$ and *ii)* $|\nabla_i f(x^{k+1})| \leq \beta$. Property *i)* follows directly by taking into account the projection (6).

Let us check property *ii)*. By using a Taylor expansion of ∇f at $x^{k+1} = x^k + \tilde{d}^k$, for some $\theta \in (0, 1)$, we get

$$\nabla f(x^{k+1}) = \nabla f(x^k) + \nabla^2 f(x^k + (1-\theta)\tilde{d}^k)^\top \tilde{d}^k. \tag{48}$$

By applying our strict complementarity hypothesis in (48) we obtain

$$|\nabla_i f(x^{k+1})| < \beta + |\sum_j \left[\nabla^2 f(x^k + (1-\theta)\tilde{d}^k)\right]_{ij} \tilde{d}_j^k|. \tag{49}$$

First, we analyze the last sum in (49). In order to simplify the notation, we define the matrix H^k with entries $H_{ij}^k = \left[\nabla^2 f(x^k + (1-\theta)\tilde{d}^k)\right]_{ij}$. By construction $\tilde{d}_i^k = 0$, and

we get the estimate

$$|\sum_{j\neq i} H_{ij}^k \tilde{d}_j^k| \leq \|H^k\|_2 \sum_{j\neq i} |\tilde{d}_j^k| \leq \|H^k\|_2 \left(\sum_{\substack{j\neq i \\ j:|x_j^k|\leq 1/\gamma}} |\tilde{d}_j^k| + \sum_{\substack{j\neq i \\ j:|x_j^k|>1/\gamma}} |\tilde{d}_j^k| \right). \tag{50}$$

For the first sum in the last inequality, each $\tilde{d}_j^k$, $j \in \mathcal{H}^k$ can be bounded using (14). For indexes $j \notin \mathcal{H}^k$ it holds that $|\tilde{d}_j^k| = |x_j^k| \leq 1/\gamma$. Therefore, there exists a constant $C > 0$ such that

$$|\sum_{j\neq i} H_{ij}^k \tilde{d}_j^k| \leq \|H^k\|_2 \left(\frac{C}{\gamma} + \sum_{\substack{j\neq i \\ j:|x_j^k|>1/\gamma}} |\tilde{d}_j^k| \right). \tag{51}$$

On the other hand, since $x^k \to \bar{x}$ as $k \to \infty$ we have that $\tilde{d}^k = x^{k+1} - x^k \to 0$ as $k \to \infty$. Consequently, for any $\varepsilon > 0$ there exists k_0 such that for any $k \geq k_0$ the relations (14) ,(49) and (51) leads to the following bound

$$|\nabla_i f(x^{k+1})| < \beta + \frac{C}{\gamma} + \varepsilon \|H^k\|_2. \tag{52}$$

Since $\varepsilon > 0$ is arbitrary and taking γ large enough, condition *ii*) holds, which together with *i*) implies (47). ∎

Remark 4. *From estimate* (52) *it becomes clear that γ plays a crucial role in order to have a larger neighbourhood where the monotonicity of the active set occurs. If γ would not be present, the index k_0, from which on the monotonicity of the strongly active sets is achieved, would be larger. This is also verified experimentally (see Section 5.2 below).*

Remark 5. *If we assume that f is separable, its Hessian is a diagonal matrix and*

$$\sum_{j\neq i} H_{ij}^k \widetilde{d}_j^k = 0,$$

which, together with the assumptions of Theorem 3, implies directly (without the strict complimentarity hypothesis) that $|\nabla_i f(x^{k+1})| = |\nabla_i f(x^k)| \leq \beta$ and $\mathcal{S}_k \subseteq \mathcal{S}_{k+1}$.

5. The reduced algorithm and its interpretation as semismooth Newton method

Based on the enriched second-order updates (9) and the subsequent projection step, the idea of combining both effects in a single update arises. A natural alternative for this consists in incorporating the projection in the building of the second order matrix. Specifically, by reordering the iterates in such a way that the components with the indexes that belong to the strongly active set $\mathcal{S}_k$ appear first, i.e., $d^k = (d_{\mathcal{S}_k}^k, d_{I\setminus\mathcal{S}_k}^k)^T$ and considering the reduced second order matrix

$$(B_R^k)_{ij} := \begin{cases} \delta_{ij} & \text{if } i \in \mathcal{S}_k, \text{ for all } j, \\ (B^k + \beta\Gamma^k)_{ij} & \text{if not,} \end{cases}$$

the following system may be solved:

$$B_R^k \begin{pmatrix} d^k_{\mathcal{S}_k} \\ d^k_{I\setminus\mathcal{S}_k} \end{pmatrix} = \begin{pmatrix} -x^k_{\mathcal{S}_k} \\ -\tilde{\nabla}\varphi(x^k)_{I\setminus\mathcal{S}_k} \end{pmatrix}. \tag{53}$$

In this manner, the second order information is only used for the update of x_i^k, $i \in I\setminus\mathcal{S}_k$. This change makes the projection step superfluous and preserves the enriched curvature information as much as possible. The complete reduced (**OESOM**) algorithm is given through the following steps:

Algorithm 2 Reduced Orthantwise Enriched Second Order Method

1: Initialize x^0 and B^0.
2: **repeat**
3: Choose the regularization parameter γ.
4: Compute the matrix Γ^k using (8) .
5: Compute the descent direction $\widetilde{\nabla}\varphi(x^k)$ using (5).
6: Compute d^k by solving the linear system (53).
7: Compute
$$x^{k+1} = x^k + d^k,$$
8: Update the matrix B^k.
9: $k \leftarrow k+1$.
10: **until** stopping criteria is satisfied

Although the theory developed for (**OESOM**) cannot be directly applied to its reduced version, we are going to show next that the reduced method can also be casted as a semismooth Newton algorithm under the choice of appropriate parameter values. Moreover, motivated by this interpretation, an adaptive regularization parameter choice strategy is built upon the reduced (**OESOM**) .

Let us start by recalling some basic notions of semismooth Newton methods and the general framework proposed in [6].

Let $D \subset \mathbb{R}^n$ be an open set. The mapping $F : D \subset \mathbb{R}^n \to \mathbb{R}^m$ is called slant (or Newton) differentiable on the open subset $V \subset D$, if there exists a generalized derivative $G : V \to \mathbb{R}^{n\times m}$ such that

$$\lim_{h\to 0} \frac{1}{\|h\|_{\mathbb{R}^n}} \|F(x+h) - F(x) - G(x+h)h\|_{\mathbb{R}^m} = 0, \tag{54}$$

for every $x \in V$. With help of this concept, a Newton iteration for finding a root $\bar{x} \in \mathbb{R}^n$ of $F(x) = 0$ is given by:

$$G(x^k)d = -F(x^k), \qquad x^{k+1} = x^k + d. \tag{55}$$

This iteration leads to a locally superlinear convergent method under suitable hypothesis on F. For more details we refer the reader to e.g. [12].

Let us now focus on our problem (1). A simple reformulation leads to the equivalent operator equation $F(x) = 0$, where

$$F_i(x) = \max\left(\min\left(\tau(\nabla_i f(x) + \beta), x_i\right), \tau(\nabla_i f(x) - \beta)\right),\ i = 1, \dots, m. \tag{56}$$

Since the *max* and *min* functions are semismooth, a Newton type update can be obtained. By defining the following index sets:

$$\begin{aligned}
\mathcal{N}^k &:= \left\{ i \colon x_i^k \leq \tau \left(\nabla_i f(x^k) - \beta \right) \right\}, \\
\mathcal{A}^k &:= \left\{ i \colon \tau \left(\nabla_i f(x^k) - \beta \right) \leq x_i^k \leq \tau \left(\nabla_i f(x^k) + \beta \right) \right\}, \\
\mathcal{P}^k &:= \left\{ i \colon x_i^k \geq \tau \left(\nabla_i f(x^k) + \beta \right) \right\},
\end{aligned}$$

the semismooth Newton updates (**SSN**) can be written in the following form:

$$e_i^T d = -x_i^k, \qquad i \in \mathcal{A}^k \setminus \left(\mathcal{N}^k \cup \mathcal{P}^k \right) \tag{57a}$$

$$\nabla^2_{i:} f(x^k) d = -\left(\nabla_i f(x^k) + \beta \right), \qquad i \in \mathcal{P}^k \setminus \mathcal{A}^k, \tag{57b}$$

$$\nabla^2_{i:} f(x^k) d = -\left(\nabla_i f(x^k) - \beta \right), \qquad i \in \mathcal{N}^k \setminus \mathcal{A}^k, \tag{57c}$$

$$\left(\delta_i \nabla^2_{i:} f(x^k) + (1 - \delta_i) e_i^T \right) d = -\tau \left(\nabla_i f(x^k) - \beta \right), \qquad i \in \mathcal{N}^k \cap \mathcal{A}^k, \tag{57d}$$

$$\left(\delta_i \nabla^2_{i:} f(x^k) + (1 - \delta_i) e_i^T \right) d = -\tau \left(\nabla_i f(x^k) + \beta \right), \qquad i \in \mathcal{P}^k \cap \mathcal{A}^k, \tag{57e}$$

$$x^{k+1} = x^k + d, \tag{57f}$$

where $\nabla^2_{i:} f(x)$ stands for the i–th row of the Hessian and e_i is the canonical vector of $\mathbb{R}^m$.

Byrd et al. [6] considered this framework for the solution of problem (**P**). By appropriately choosing the algorithm parameters, several well-known algorithms (**FISTA** [3], **OWL** [2]) may be derived from this rather general setting. The orthant-wise method proposed by Byrd et al. [6], for instance, is equivalent to the semismooth Newton updates under the special choice of τ sufficiently small and

$$\delta_i = 0, \qquad \text{for all } i \in \left(\mathcal{N}^k \cap \mathcal{A}^k \right) \cup \left(\mathcal{P}^k \cap \mathcal{A}^k \right).$$

In the case of our reduced algorithm, by choosing $\tau = 1/(\beta\gamma+1)$ and γ sufficiently large, the equivalence with semismooth Newton updates is obtained.

Theorem 4. *For γ sufficiently large and $\tau = 1/(\beta\gamma+1)$, the iterations generated by Algorithm 2 are equivalent to the semismooth Newton iterations obtained through the solution of* (57).

Proof. By choosing $\tau = 1/(\beta\gamma+1)$ and γ sufficiently large we obtain that

$$\operatorname{sign}\left(x_i^k - \tau \left(\nabla_i f(x^k) + \operatorname{sign}(x_i^k)\beta \right) \right) = \operatorname{sign}(x_i^k) \quad \text{for all } i : x_i^k \neq 0. \tag{58}$$

The latter implies that if $x_i^k \in \mathcal{A}^k$, then $x_i^k = 0$. Indeed, if $\operatorname{sign}(x_i^k) = -1$ then (58) implies that $x_i^k - \tau \left(\nabla_i f(x^k) - \beta \right) < 0$ and, consequently, $i \in \mathcal{N}^k \setminus \mathcal{A}^k$. In a similar way, if $x_i^k > 0$ then $i \in \mathcal{P}^k \setminus \mathcal{A}^k$.

Concerning the updates:

- If $i \in \mathcal{A}^k \setminus \left(\mathcal{N}^k \cup \mathcal{P}^k \right)$, then by (57a)

$$d_i^k = -x_i^k = 0.$$

On the other hand, in our method this case corresponds to $z_i^k = 0$, which thanks to the incorporation of the projection in the building of the reduced second order matrix implies that

$$x_i^{k+1} = x_i^k = 0.$$

- If $\operatorname{sign}(x_i^k) = -1$, then $i \in \mathcal{N}^k \setminus \mathcal{A}^k$ and the (**SSN**) update corresponds to

$$\nabla^2_{i:} f(x^k) d^k = -\left(\nabla_i f(x^k) + \beta\right),$$

which is similar to the update of the reduced (**OESOM**) algorithm.

- If $\operatorname{sign}(x_i^k) = 1$, then $i \in \mathcal{P}^k \setminus \mathcal{A}^k$ and

$$\nabla^2_{i:} f(x^k) d^k = -\left(\nabla_i f(x^k) - \beta\right),$$

which corresponds to the update of our reduced algorithm.

- If $i \in \mathcal{N}^k \cap \mathcal{A}^k$, then our algorithm yields the update

$$\left(\nabla^2_{i:} f(x^k) + \beta\gamma e_i^T\right) d^k = -\left(\nabla_i f(x^k) - \beta\right), \tag{59}$$

which corresponds exactly to the (**SSN**) update (57d) with the choice $\delta = {}^1\!/\!_{(\beta\gamma+1)}$. Indeed,

$$\left(\delta_i \nabla^2_{i:} f(x^k) + (1-\delta_i) e_i^T\right) d = -\tau\left(\nabla_i f(x^k) - \beta\right)$$

$$\Leftrightarrow \left(\frac{1}{\beta\gamma+1}\nabla^2_{i:} f(x^k) + \frac{\beta\gamma}{\beta\gamma+1} e_i^T\right) d = -\frac{1}{\beta\gamma+1}\left(\nabla_i f(x^k) - \beta\right),$$

which coincides with (59). The case $i \in \mathcal{P}^k \cap \mathcal{A}^k$ follows in a similiar way.

■

Consequently, the updates of the reduced (**OESOM**) and the (**SSN**) method coincide for the special choice $\tau = \delta = \dfrac{1}{\beta\gamma+1}$. Let us remark that for this to occur, full steps ($s = 1$) have to be performed.

From the semismooth interpretation also an adaptive strategy for the regularization parameter can be divised. Indeed, we may choose τ such that

$$\operatorname{sign}\left(x_i^k - \tau\left(\nabla_i f(x^k) + \operatorname{sign}(x_i^k)\beta\right)\right) = \operatorname{sign}(x_i^k), \quad \forall i\colon x_i^k \neq 0$$

is satisfied. This leads to the condition

$$\tau < \frac{|x_i^k|}{|\nabla_i f(x^k) + \operatorname{sign}(x_i^k)\beta|}, \qquad \forall i\colon x_i^k \neq 0.$$

Considering the choice $\tau = \delta = {}^1\!/\!_{(\beta\gamma+1)}$ of our method, an adaptive choice of γ is given by

$$\gamma > \frac{|\nabla_i f(x^k) + \operatorname{sign}(x_i^k)\beta|}{\beta|x_i^k|} - \frac{1}{\beta}, \qquad \forall i\colon x_i^k \neq 0,$$

or, more conservatively,

$$\gamma = \max_{\{i\colon x_i^k \neq 0\}} \left(\frac{|\nabla_i f(x^k) + \operatorname{sign}(x_i^k)\beta|}{\beta|x_i^k|}\right) \tag{60}$$

The adaptive reduced (**OESOM**) algorithm is obtained by using equation (60) in step 3 of Algorithm 2.

6. Numerical experiments

The computational study of our algorithm is divided in two sets of numerical experiments. The first set of problems is intended to compare (**OESOM**) with other state–of-the–art methods for solving ℓ_1 penalized problems. Specifically, we compare our method with other second order methods as the Orthant–wise Limited–memory Quasi–Newton method (**OWL**) [2], Newton–CG algorithm (**NW-CG**) [6], primal–dual Newton–CG (**pdNCG**) [13]. In addition, we compare our method with the popular first–order algorithm Fast Iterative Shrinkage–thresholding Algorithm (**FISTA**) which is an accelerated proximal method (see, e.g., [3]).

The second set of experiments focuses on the numerical properties of (**OESOM**). Monotonicity properties are investigated by means of the strong active sets defined by z^k. We also consider a numerical continuation strategy for the solution of the problem with different levels of sparsity associated to the parameter β, as this is frequently used.

The experiments show that (**OESOM**) is competitive compared to other state–of–the–art methods, and in many cases is able to improve convergence in terms of number of iterations and/or execution time. We emphasize that (**OWL**) and (**NW-CG**) use second–order information from the regular part only. In the case of (**pdNCG**), a second order regularization of the ℓ_1–norm is used from the beginning and, consequently, a regularized version of the original problem is actually solved. In our algorithm the second order matrix of the regular part is modified by Γ^k (see (9)), which has an scaling effect on the directions associated with those components that are potentially active. We observe that this feature boosts the active set identification process.

The most computationally expensive step of our algorithm is the solution of the system (9). We address this issue by applying Krylov methods, such as Arnoldi [29], which exploit the structure of the matrix $B^k + \beta\Gamma^k$ and lead to a fast iterative solution of the system. This is studied in depth in Subsection 6.2, where an inexact version of (**OESOM**) is also considered. Section 5 is devoted to the numerical testing of the reduced version of (**OESOM**) developed in Section 5 which considerably reduces the size of the associated linear system (9).

Implementation aspects. The code for (**OWL**) was implemented in MATLAB according to the original paper [2]. The (**NW-CG**) algorithm was also implemented in MATLAB according to the original paper [4]. In the case of (**FISTA**), we have used the code available in the *Toolbox of Sparse Optimization* from the Matlab File Exchange repository. The MATLAB code for (**pdNCG**), was publicly provided by the authors (`http://www.maths.ed.ac.uk/ERGO/pdNCG/`). Unless otherwise stated, in our experiments we have used a pseudo–Huber regularization parameter $\mu = 1e-7$.

For measuring the numerical efficiency of the different algorithms, we have considered both the number of iterations and the execution time. In particular, the number of iterations is relevant in the context of PDE-constrained optimization, where at every step two partial differential equations (state and adjoint equations) must be solved to compute descent directions and evaluate the objective function. For the stopping criterion, in general, we use

$$\|\widetilde{\nabla}\varphi(x)\|_\infty < \eta,$$

for a given tolerance $0 < \eta \ll 1$. However, we use other suitable criteria when comparing with other algorithms, since the provided codes for other algorithms do not necessarily use the same stopping criteria. In each experiment, if necessary, we specify how the stopping criteria, the parameters and function values are chosen.

		Number of iterations									
Size of A		OESOM		NW–CG		OWL		pdNCG		FISTA	
m	n	MEAN	SDV	MEAN	SDV	MEAN	SDV	MEAN	SDV	MEAN	SDV
400	200	8.20	1.475	271.60	30.685	15.40	4.948	13.90	1.595	1000	105.4093
800	400	8.60	0.966	291.40	38.982	16.80	3.293	14.20	1.316	1430	122.927
1200	600	8.80	1.135	193.50	162.710	19.70	7.409	14.30	0.674	1902	127.021
1600	800	9.70	0.948	235.30	128.154	17.70	3.093	15	0.667	2300	246.080
2000	1000	11.30	3.772	271.50	98.592	17.80	2.658	13.60	0.699	2883	193.278
2400	1200	14.90	5.782	304.50	49.996	21.40	6.963	14.30	1.337	3085	158.201

TABLE 1. Comparison of mean and standard deviation of the number of iterations for the random LASSO problems

		Time (s)									
Size of A		OESOM		NW–CG		OWL		pdNCG		FISTA	
m	n	MEAN	SDV	MEAN	SDV	MEAN	SDV	MEAN	SDV	MEAN	SDV
400	200	0.0881	0.0576	1.5984	0.4235	0.1978	0.1236	2.5302	0.5611	0.2908	0.0371
800	400	0.1825	0.0251	6.0516	0.9539	0.3954	0.0641	12.5289	2.5771	1.3778	0.1496
1200	600	0.4485	0.0560	7.7453	6.7164	0.9660	0.1612	54.0916	14.0477	2.4756	0.2478
1600	800	1.0372	0.0897	20.6573	11.7550	2.4216	0.3841	130.9736	17.9134	3.2964	0.4936
2000	1000	2.5998	0.9613	51.0575	19.1155	4.7951	0.7360	305.1327	95.1593	7.4477	1.1067
2400	1200	5.1059	1.9218	85.0013	18.3755	7.2795	0.7749	442.3785	121.0738	15.0734	3.9490

TABLE 2. Comparison of mean and standard deviation of the execution time for the random LASSO problems

We mention that for the numerical approximation of the partial differential equations involved, a standard finite difference discretization scheme was used, where the Laplacian was approximated using the five points stencil [28].

6.1. **Numerical comparison with other state–of–the–art algorithms.**

6.1.1. *Randomly generated linear-quadratic optimization problems.* For our first experiment, we consider the classical least–squares sparse problems (LASSO). These problems are randomly generated according to the procedure developed in [25, Section 6]. LASSO problems have the following form:

$$\min_{x\in\mathbb{R}^n} \varphi(x) = \frac{1}{2} \parallel Ax - b \parallel_2^2 + \|x\|_1,$$

where $A \in \mathbb{R}^{m\times n}$ and $b \in \mathbb{R}^m$, with $m \geq n$ to assure that the problem has a unique solution. In this procedure, we select $m^* < m$ as number of sparse components in the solution x^*, and its objective function value φ^*. For more details on the construction and generation of these problems, we refer to [25, Section 6].

We consider 6 sets of problems of different sizes. For each size we generate 10 random problems to compare with the algorithms mentioned at the beginning of this section. Since the third–party codes use different stopping criteria, and we cannot always guaranteed to satisfy an exact cost value, we have run the algorithms until:

$$|\varphi^* - \varphi(x^k)| \leq 10^{-5},$$

In our experiments, all the algorithms reach the stoping criteria with an error $\|x^k - x^*\|_2$ of order 10^{-4}. Table 1 sums up our numerical findings for each experiment where the mean and the standard deviation of the number of iterations and execution time, show that (**OESOM**) is able to outperform the others. For the OWL algoritm we fix to 20 the number of vectors for the L–BFGS method.

Algorithm	Number of Iterations	Time (s)	Cost function
OESOM	10	2.94	1.5636
NW–CG	347	37.18	1.5636
OWL	19	10.90	1.5636
FISTA	700	155.68	1.5636

TABLE 3. Comparison of execution time and number of iterations for PDE-constrained optimization problems.

6.1.2. *PDE-constrained optimization.* Let us now consider the following linear-quadratic optimal control problem:

$$\text{(OCP)} \qquad \begin{cases} \min\limits_{(y,u)} \ \dfrac{1}{2}\|y - y_d\|^2_{L^2(D)} + \dfrac{\alpha}{2}\|u\|^2_{L^2(D)} + \beta\|u\|_{L^1(D)} \\ \text{subject to} \\ \qquad -\nu\Delta y = u \quad \text{in } D := (0,1)\times(0,1), \\ \qquad\quad y = 0 \qquad \text{on } \partial D. \end{cases}$$

The numerical solution of this type of problems is an interesting experiment for testing our algorithm since these kind of problems involve a high computational cost in every iteration, due to the presence of the PDE–constraints which requires the numerical solution of partial differential equations in order to compute the gradient of f in each iterate.

For this problem, we consider the following parameter values $\nu = 1$, $\alpha = 2e-5$, $\beta = 9.4e-4$ and $y_d := \sin(4\pi x)\cos(8\pi y)\exp(2\pi x)$. We use a *discretize-then-optimize* approach, transforming the original infinite-dimensional optimization problem into a finite-dimensional one in the form of (**P**). This is done by means of the finite difference method. In the following experiments we have discretized the unit square using a uniform grid with 3600 internal nodes.

We compare the performance of (**OESOM**) with respect to other three different methods. For the purpose of comparison, we have considered a target value of the cost function as a stopping criterion:

$$\varphi(x^k) < 1.5637, \tag{61}$$

and measure the execution time and number of iterations that each of the algorithms needs to reach the prescribed cost value. The obtained numerical results are shown in Table 3, where we observe that (**OESOM**) requires fewer iterations to converge.

Figure 1 illustrates the active set evolution of (**OESOM**) , (**OWL**) and (**NW–CG**) along their iterations. The null components of z^k are depicted in Figure 1a and the decay of the objective function is shown in Figure 1b. It can be noticed that (**OESOM**) seems to be faster at identifying the active sets. This feature plays a significant role in computing the optimal control. In addition, in several cases we observe that (**OESOM**) algorithm attains a smaller value of the objective function as is shown in Table 4.

The numerical results for this type of PDE-constrained optimization problems are strongly affected by the choice of the diffusion parameter ν and the weight β. For small values of ν and β close to its critical value β_0, the solution becomes more sparse and harder to obtain. To evaluate the performance for such cases, we test the algorithms (**OWL**) , (**NW-CG**) and (**OESOM**) for different combinations of $(\nu, \beta) \in [0.0743, 0.6] \times [0.0120, 0.02]$ on a grid of 625 points and plot the corresponding number of iterations in Figure 2. Dark red stands for a high number of iterations (100)

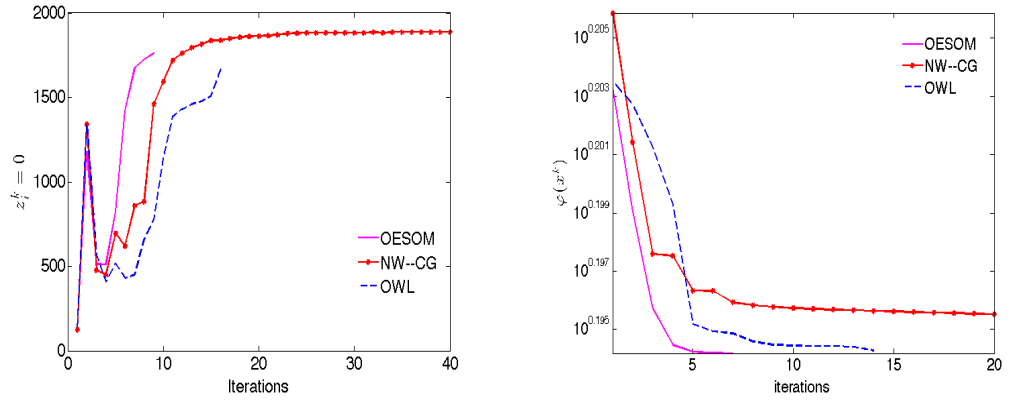

(A) Null component of z^k. (B) Cost function.

FIGURE 1. Performance for PDE-constrained optimization problems

and blue for a small number according to the color scale. It can be observed that for this set of experiments, in general (**OESOM**) needs less iterates than the other two to reach the solution and exhibits a robust behaviour with respect to the parameters. Some of the results are also presented in Table 4, including (**FISTA**).

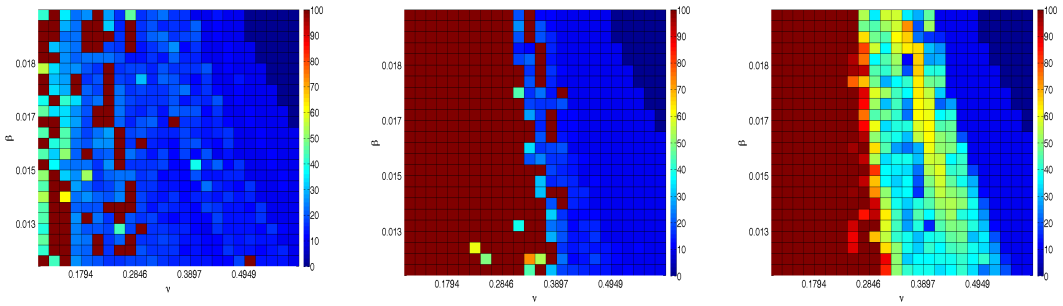

FIGURE 2. Number of iterations: OESOM (left), NW–CG (center), OWL (right). Dark red is for high number for iterations and blue for low ones.

Regularization Parameters		ITERATIONS				TIME (s)				COST FUNCTION			
		OE	NC	OW	FT	OE	NC	OW	FT	OE	NC	OW	FT
α	1e-5	8	100	21	470	3.08	31.40	7.47	150.62	1.5263	1.5376	1.5267	1.5263
β	0.0012												
α	1.2e-5	8	100	25	462	2.97	29.05	7.43	142.87	1.5515	1.5564	1.5524	1.5515
β	0.0014												
α	1.4e-5	8	100	23	470	2.92	29.79	7.39	160.23	1.5695	1.5763	1.5700	1.5696
β	0.0016												
α	3e-5	9	100	11	470	3.87	14.17	7.06	154.26	1.6149	1.6154	1.6150	1.6149
β	0.0025												

TABLE 4. Comparison with different regularization parameters for PDE–constrained problems. OE stands for OESOM, NC for NW–CG, OW for OWL and FT for FISTA.

In Table 5 the results of this experiment for different values of the Huber regularization parameter γ are registered. The behaviour of the algorithm appears to be robust with respect to the parameter.

γ	Iterations	Execution Time (s)	Cost Function
1e3	13	4.04	1.5642
1e4	8	2.11	1.5641
1e5	14	4.64	1.5647

TABLE 5. Performance of (**OESOM**) for different values of γ.

6.1.3. *Machine learning: dataset training problem.* Many problems arising in machine learning involve a training step (for a given dataset) in order to determine the parameters of a certain model for data classification or feature selection (see, e.g., [35]). This training step consists in solving an optimization problem of the form (**P**):

$$\min_{x\in\mathbb{R}^m} \ell(x) + \beta\|x\|_1. \tag{62}$$

Such problems typically involve a multi-class logistic function ℓ, known as loss function, that represents the normalized sum of the negative log likelihood of each data point being placed in the correct class [10, 24] and it is defined by

$$\ell(x) = -\frac{1}{N}\sum_{j=1}^{N} \log \frac{\exp(x_{y_j}^T z_j)}{\sum_{i\in C}\exp(x_i^T z_j)},$$

where N is the number of samples used for the recognition, C denotes the set of all class labels, y_j the label associated to the training points j, z_j is the feature vector and x_i is the parameters' subvector of class label i. Again, the parameter β afects the sparsity of the solution as it increases. In this context the non–zero entries of the solution are interpreted to be the most representative parameters for the classification function.

In our experiment we train a Statlog–Satellite database, consisting of sub-area images of size 82 x 100 pixels. The aim is to classify pixels from satellite images as *red soil*, *cotton crop*, *damp grey soil*, *soil with vegetation stubble*, *mixture class* and *very damp grey soil*. More details of this database can be found in [19].

We test the three algorithms (**OESOM**) , (**NW-CG**) and (**OWL**) for solving this problem. The results are shown in Table 6. It can be observed that our method slightly outperforms the others, with respect to execution time, for this particular experiment.

Algorithm	**Cost function**	**Time (h)**
OESOM	0.9361	1.14
NW–CG	0.9361	1.63
OWL	0.9396	1.58

TABLE 6. Time spent in training a dataset training problem

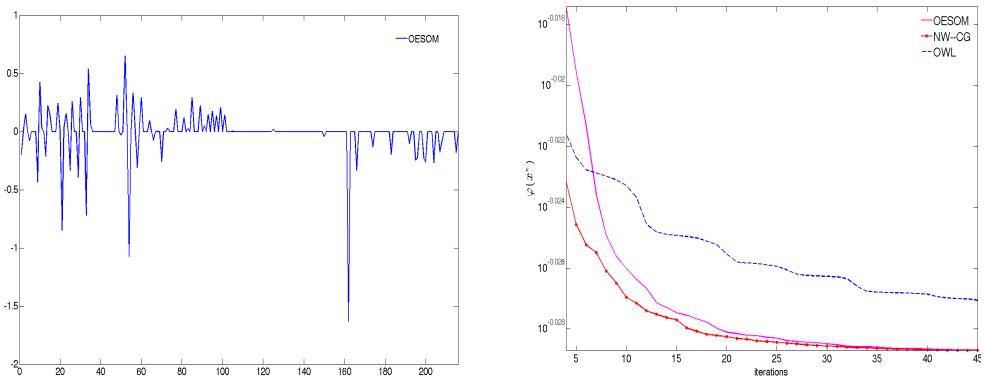

FIGURE 3. Training data set example: Solution (left) and evolution of the cost function value (right)

6.2. **Numerical properties of (OESOM).**

6.2.1. *Monotonicity of active sets.* On basis of the analysis carried out in Section 4, we expect a monotone behavior of the null components of z^k (strong active set) in a neighbourhood of the solution. Furthermore, γ theoretically affects the size of such neighbourhood in the case of (**OESOM**) . To test this, we consider first the PDE–constrained optimization experiments from Subsection 6.1.2. In Figure 4 we show the size of the null components of the orthant direction z^k in each iterations, for several choices of the regularization parameters α and β. We observe that (**OESOM**) exhibits a monotone increase in the cardinality of the active set earlier than the other algorithms.

Next we consider a similar monotonicity test for randomly generated LASSO problems as described in Subsection 6.1.1. We consider four different sizes of problems and take 10 samples in each case. In Figure 5 we color a square with blue in case a larger active set is reached in that iteration with respect to the previous one, otherwise we color that square with red. A large dominant blue behaviour is observed in Figure 5, where the red squares are rarely present. This is in agreement with our theoretical findings, mainly Theorem 4.

6.2.2. *Reduced Orthantwise Enriched Second Order Method (R–OESOM).* In Section 5, we introduced Algorithm 2 which may be interpreted as a semismooth Newton method by an appropriate choice of the regularization parameter γ. One important advantage of this algorithm is that size of the linear system (53) is considerably smaller than (9), depending to how sparse the solution is. We next present the numerical perfromance of this reduction strategy by solving the set of LASSO problems described in Section 6.1.1, and comparing with the original (**OESOM**) . The numerical results of this comparison are presented in Table 7, where a competitive bahaviour of the reduced method can be observed in terms of execution time and number of iterations.

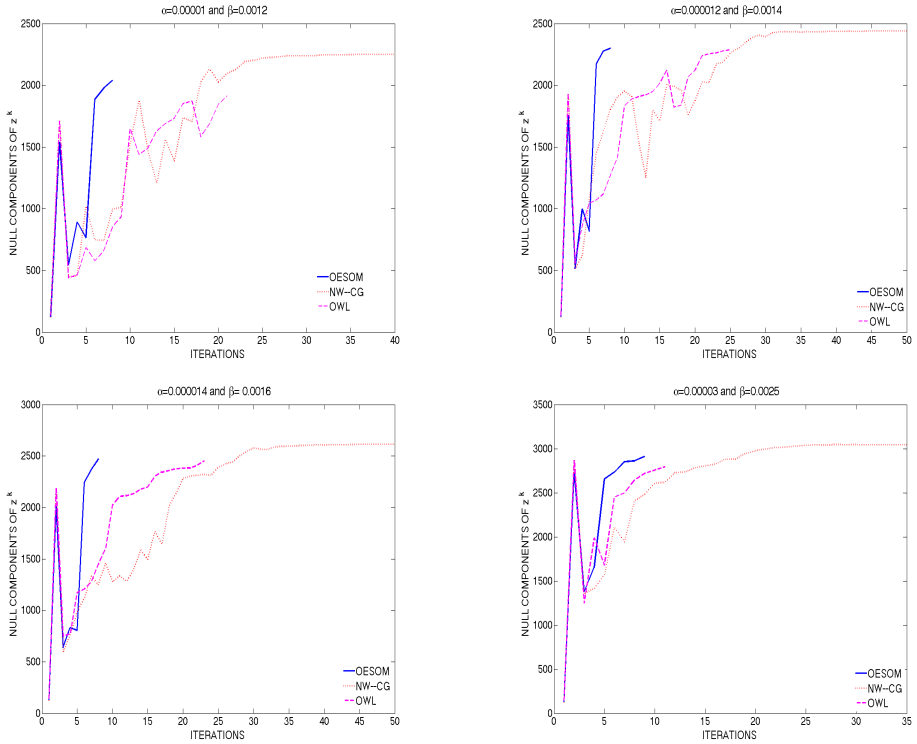

FIGURE 4. Evolution of the active sets: PDE-constrained optimization test.

		Time (s)				Iterations			
SIZE		OESOM		REDUCED OESOM		OESOM		REDUCED OESOM	
m	n	MEAN	SDV	MEAN	SDV	MEAN	SDV	MEAN	SDV
400	200	0.0881	0.0576	0.0810	0.0431	8.20	1.475	8.10	1.5239
800	400	0.1825	0.0251	0.1600	0.0206	8.60	0.9661	8.20	1.2293
1200	600	0.4486	0.0560	0.4616	0.1062	8.80	1.1353	8.20	1.2293
1600	800	1.0372	0.0897	0.7417	0.0592	9.70	0.9487	7.60	0.5164
2000	1000	2.5998	0.9613	1.7928	0.2843	11.30	3.7727	7.80	0.9189
2400	1200	5.1059	1.9218	2.1667	0.2780	14.90	5.7822	7.50	0.5270

TABLE 7. Comparison of (**OESOM**) and **(R–OESOM)** for randomly generated LASSO problems.

6.2.3. *Inexact Orthantwise Enriched Second Order Method (I–OESOM).* In order to make (**OESOM**) even more efficient, we consider an inexact variant of the algorithm, where the associated linear system is solved only approximately in each iteration, according to the following rule:

$$\| (B_k + \Gamma_k)\, d^k + \widetilde{\nabla}\varphi_k \| \leq \xi \, \| \widetilde{\nabla}\varphi_k \|, \tag{63}$$

where ξ is a chosen tolerance. The linear system is solved (inexactly) by using Arnoldi's method (see, e.g., [30]).

We use the PDE-constrained optimization problem of Section 6.1.2 to test the inexact variant of (**OESOM**) on the discretized unit square with 3844 internal nodes. Moreover, we have set the parameters $\alpha = 2e-5$, $\nu = 1$ and $\gamma = 1e4$. In Table 8 we compare the performance with respect to the original (**OESOM**) algorithm and different tolerances ξ for the inexact strategy. The last two rows of the table correspond to

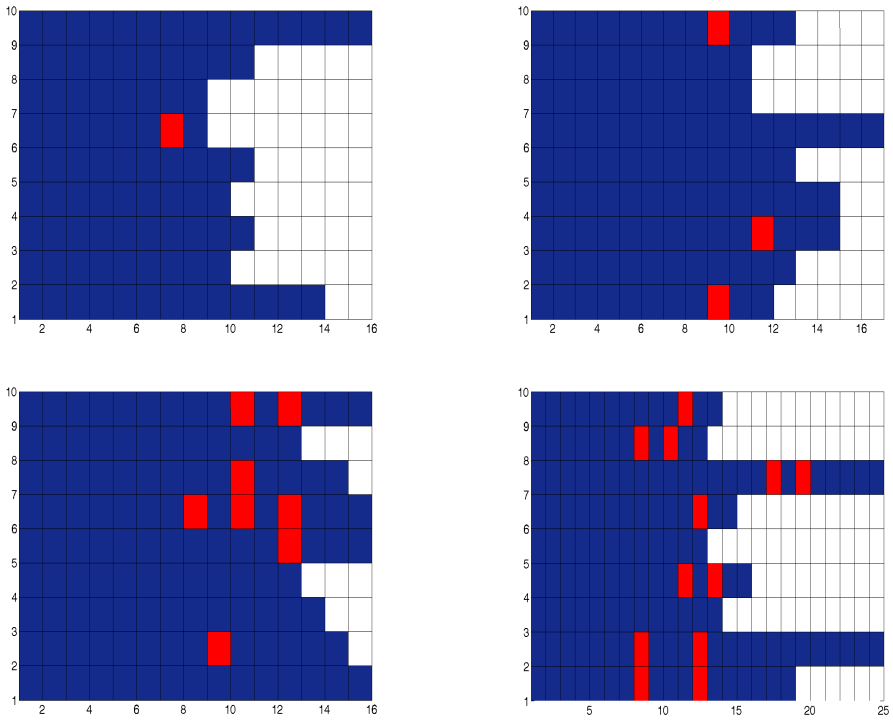

FIGURE 5. Evolution of the active sets: Randomly generated LASSO. Size of the problems: Upper-left (400×200), upper-right (800×400),lower-left (1200×600) and lower-right (1600×800).

variable tolerances defined by $\xi_k = \{(1/2)^k\}$ and $\xi_k = \| \widetilde{\nabla}\varphi(x_k) \|$ commonly used in the literature (see, e.g., [15]). From these results we may infer that the inexact strategy actually helps to reduce the computational cost of the corresponding (**OESOM**) variant, without damaging the convergence properties of the method. This fact actually deserves future theoretical investigation.

Algorithm	**Iterations**	**Cost function**	**Time (s)**
OESOM	9	1.564	9.46
I–OESOM ($\xi = 1e-1$)	9	1.564	2.58
I–OESOM ($\xi = 1e-2$)	8	1.564	2.8
I–OESOM ($\xi = 1e-3$)	8	1.564	3.02
I–OESOM ($\xi = (1/2)^k$)	8	1.564	2.15
I–OESOM ($\xi = \| \widetilde{\nabla}\varphi_k \|$)	12	1.565	3.86

TABLE 8. Numerical performance of the inexact version of (**OESOM**) .

6.2.4. *Varying the sparsity penalization coefficient β.* In this experiment we consider again the optimal control problem (OCP) with $\alpha = 0.00002$. Here, we reconstruct the solution path for different sparsity levels by changing the value of β up to the critical value $\beta_0 = \| \widetilde{\nabla}\varphi(0) \|_\infty$. Figure 6 shows plots a logarithmic growth for the sparsity of the solution as well as its associated cost.

β	Iterations		Time (s)	
	WC	C	WC	C
0.009	8	8	11.0625	11.0650
0.0019	12	5	19.3686	6.6957
0.0030	8	5	10.7736	6.7271
0.0040	11	4	17.1299	4.2611
0.0050	10	5	15.2321	6.7587
0.0060	13	3	21.9843	2.2405
0.0070	13	5	23.8750	6.2499
0.0080	11	3	19.7571	2.2646
0.0090	2	5	0.0479	0.1183
0.0100	2	1	0.0467	0.0207

TABLE 9. Numerical performance of (**OESOM**) for different levels of sparsity. WC stands for the numerical performance without continuation, and C with continuation.

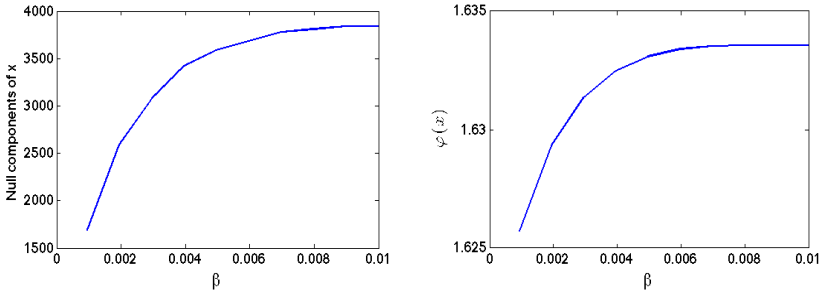

FIGURE 6. Null components of solution (left) and cost function value (right) for different values of β

In order to measure the efficiency of (**OESOM**) in reconstructing the solution path, we evaluate its performance with and without a continuation strategy. First, we solve this family of problems with the same initial iterate for each different value of β. Then we solve the same family of problems by a continuation strategy, where the optimal solution for the previous β value is used to initialize the algorithm to solve the problem with the next β value. As expected, the continuation strategy is considerably more efficient, as presented in Table 9.

7. CONCLUSIONS

By using weak information of the ℓ_1-norm, through partial Huber regularization, we are able to obtain extra second-order information of the objective function, which is integrated in the computation of the descent direction. This information turns out to be important for a faster identification of the active-sets and improved convergence properties of the resulting algorithm.

A reduced version of the proposed method has been proved to be equivalent to a semismooth Newton scheme with the special choice $\tau = \delta = \frac{1}{\gamma+1}$ of the (**SSN**) parameters. Since the (**SSN**) is a fast local method, such convergence properties are also inherited by our algorithm. Moreover, thanks to this interpretation, an adaptive update strategy for the regularization parameter has been proposed.

Finally, the performance of the proposed algorithms turns out to be competitive with respect to other state-of-the-art methods. Specifically, we exhaustively compared the efficiency of (**OESOM**) with respect to the second–order algorithms (**OWL**) ,

(**NW-CG**) and (**pdNCG**), and the popular first-order algorithm (**FISTA**), for three different families of large-scale optimization problems. Also an inexact variant of (**OE-SOM**) was tested with promising results. From the numerical experiments carried out, a competitive performance of the proposed algorithms was verified.

References

[1] Libsvm data: Classification (binary class). `http://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/binary.html#heart`. Accessed: 2016-05-24.

[2] G. Andrew and J. Gao. Scalable training of $\ell_1$regularized log-linear models. *In Proceedings of the Twenty Fourth Conference on Machine Learning (ICML)*, 2007.

[3] A. Beck and M. Teboulle. A Fast Iterative Shrinkage-Thresholding Algorithm for Linear Inverse Problems. *SIAM Journal on Imaging Sciences*, 2(1):183–202, March 2009.

[4] R. Byrd, G. Chin, J. Nocedal, and Y. Wu. Sample size selection in optimization methods for machine learning. *Mathematical Programming*, 134(1), 2011.

[5] R. Byrd, G.M. Chin, W. Neveitt, and J. Nocedal. On the use of stochastic Hessian information in unconstrained optimization. *SIAM J. Optim*, 21(3):977–995, 2011.

[6] R. Byrd, G.M. Chin, J. Nocedal, and F. Oztoprak. A family of second-order methods for convex $\ell_1$regularized optimization. *Mathematical Programming*, pages 1–33, 2012.

[7] E. Casas, C. Ryll, and F. Tröltzsch. Sparse optimal control of the Schlögl and Fitzhugh–Nagumo systems. *Computational Methods in Applied Mathematics*, 13(4):415–442, 2013.

[8] E. Chouzenoux, J.C. Pesquet, and A. Repetti. Variable metric forward–backward algorithm for minimizing the sum of a differentiable function and a convex function. *Journal of Optimization Theory and Applications*, 162(1):107–132, 2014.

[9] P. Ciarlet. *Linear and nonlinear functional analysis with applications*. SIAM, 2013.

[10] M. Collins and T. Koo. Discriminative reranking for natural language parsing. *Computational Linguistics*, 31(1):25–70, 2005.

[11] J. N Darroch and D. Ratcliff. Generalized iterative scaling for log-linear models. *The annals of mathematical statistics*, pages 1470–1480, 1972.

[12] F. Facchinei and J.S. Pang. *Finite-dimensional Variational Inequalities and Complementarity Problems, Vols. I and II*. Springer, Berlin, 2003.

[13] Gonzio J. Fountoulakis, K. A second-order method for strongly convex ℓ_1-regularization problems. *Mathematical Programming*, 156(1):189–219, 2016.

[14] H. Fu, M.K. Ng, M. Nikolova, and J.L. Barlow. Efficient minimization methods of mixed ℓ_2-ℓ_1 and ℓ_1-ℓ_1 norms for image restoration. *SIAM J. Sci. Comput.*, 27(6):1881–1902, 2006.

[15] C. Geiger and C. Kanzow. *Numerische Verfahren zur Lösung unrestringierter Optimierungsaufgaben*. Springer-Lehrbuch. Springer, 1999.

[16] G.H. Golub and C.F. Van Loan. *Matrix computations*, volume 3. JHU Press, 2012.

[17] R. Herzog, G. Stadler, and G. Wachsmuth. Directional sparsity in optimal control of partial differential equations. *SIAM Journal on Control and Optimization*, 50(2):943–963, 2012.

[18] S-I. Lee, H. Lee, P. Abbeel, and A.Y. Ng. Efficient ℓ_1 regularized logistic regression. In *Proceedings of the National Conference on Artificial Intelligence*, volume 21, page 401. Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999, 2006.

[19] M. Lichman. UCI machine learning repository, 2013.

[20] R. Malouf. A comparison of algorithms for maximum entropy parameter estimation. In *proceedings of the 6th conference on Natural language learning-Volume 20*, pages 1–7. Association for Computational Linguistics, 2002.

[21] OL. Mangasarian and WH. Wolberg. Cancer diagnosis via linear programming. university of wisconsin-madison. *Computer Sciences Department*, 1990.

[22] N. Meinshausen and B. Yu. Lasso-type recovery of sparse representations for high-dimensional data. *The Annals of Statistics*, pages 246–270, 2009.

[23] A. Milzarek and M. Ulbrich. A semismooth newton method with multidimensional filter globalization for ℓ_1-optimization. *SIAM Journal on Optimization*, 24(1):298–333, 2014.

[24] T.P. Minka. A comparison of numerical optimizers for logistic regression. *Unpublished draft*, 2003.

[25] Y. Nesterov. Gradient methods for minimizing composite functions. *Mathematical Programming*, 140(1):125–161, 2013.

[26] J. Nocedal and S.J. Wright. *Numerical Optimization*. Springer-Verlag, New York, 1999.

[27] MJD. Powell. On the convergence of the variable metric algorithm. *IMA Journal of Applied Mathematics*, 7(1):21–36, 1971.
[28] A. Quarteroni. *Numerical models for differential problems*, volume 2. Springer Science & Business Media, 2010.
[29] Y. Saad. Krylov subspace methods for solving large unsymmetric linear systems. *Mathematics of computation*, 37(155):105–126, 1981.
[30] Y. Saad. *Iterative methods for sparse linear systems*. Siam, 2003.
[31] Marianna De Santis, Stefano Lucidi, and Francesco Rinaldi. A fast active set block coordinate descent algorithm for l_1-regularized least squares. *SIAM Journal on Optimization*, 26(1):781–809, 2016.
[32] M. Schmidt, G. Fung, and R. Rosales. Fast optimization methods for ℓ_1 regularization: A comparative study and two new approaches. In *Machine Learning: ECML 2007*, pages 286–297. Springer, 2007.
[33] F. Sha and F. Pereira. Shallow parsing with conditional random fields. In *Proceedings of the 2003 Conference of the North American Chapter of the Association for Computational Linguistics on Human Language Technology-Volume 1*, pages 134–141. Association for Computational Linguistics, 2003.
[34] Stefan Solntsev, Jorge Nocedal, and Richard H Byrd. An algorithm for quadratic 1-regularized optimization with a flexible active-set strategy. *Optimization Methods and Software*, 30(6):1213–1237, 2015.
[35] S. Sra, S. Nowozin, and S.J. Wright. *Optimization for machine learning*. MIT Press, 2012.
[36] G. Stadler. Elliptic optimal control problems with L^1-control cost and applications for the placement of control devices. *Comput. Optim. Appl.*, 44(2):159–181, 2009.
[37] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 267–288, 1996.
[38] F. Tröltzsch. *Optimal control of partial differential equations*, volume 112 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2010. Theory, methods and applications, Translated from the 2005 German original by Jürgen Sprekels.
[39] J.H. Wilkinson. *The algebraic eigenvalue problem*, volume 87. Clarendon Press Oxford, 1965.
[40] Stephen J Wright. Accelerated block-coordinate relaxation for regularized optimization. *SIAM Journal on Optimization*, 22(1):159–186, 2012.
[41] G. Yuan, K. Chang, C. Hsieh, and C. Lin. A comparison of optimization methods and software for large-scale ℓ_1-regularized classification. *Journal of Machine Learning Research*, (11):3183–3234, 2010.

‡Research Center of Mathematical Modelling (MODEMAT) and Department of Mathematics, Escuela Politécnica Nacional, Quito, Ecuador